# The Death of a Search Engine

## The Progression of Content Discovery, the Mechanics of Intent, and What the Answer-First Web Asks of Businesses

Himanshu Bisht
Synthesizer Lab LLC
July 2026

### Abstract

The first point of contact between a person and the internet has changed shape five times in three decades: from human-curated directories to crawler-based keyword engines, to link-authority ranking, to enriched result pages, and now to answer-first interfaces built on transformer language models. This paper traces that progression and examines what it does to user behavior and to the economics of being found. Drawing on clickstream panels, click-through studies, and the technical lineage that runs from the attention mechanism of Vaswani et al. (2017) through retrieval-augmented generation, the paper formalizes discovery as an expected-visit model in which traffic decomposes into a shrinking ranked-list term and a growing citation term. The evidence reviewed shows a consistent pattern: sixty-eight percent of US Google searches now end without a click, click-through to top-ranked pages falls by more than half when an AI summary is present, and yet visitors who do arrive from AI systems convert at materially higher rates. The paper then revisits Broder's intent taxonomy under synthesis, showing which intent classes the answer layer absorbs and which it leaves intact, and derives five imperatives for businesses: be retrievable, be citable, be corroborated, be measurable, and be worth the click. The strategic variable has moved from position on a page to probability of inclusion in an answer.

**Keywords:** content discovery; search engines; answer engines; user intent; zero-click search; generative engine optimization; retrieval-augmented generation; consumer search behavior

**JEL Classification:** D83, L86, M31, M15

## 1. Introduction

For most people, the internet begins with a question. Where that question goes, and what happens in the seconds after it is asked, has been remade five times in thirty years. A student in 1994 browsed a hand-built directory of categorized links. Her counterpart in 2004 typed keywords into a search box and scanned ten blue links. By 2014 the results page had started answering simple questions on its own, and by 2024 a growing share of questions never produced a results page at all. They produced a paragraph: composed on demand by a language model, grounded in retrieved sources, and often consumed without a single click.

The behavioral evidence behind that last turn is hard to dismiss. In the first four months of 2026, 68 percent of US Google searches ended without a click to the open web, up from 58.5 percent in 2024 (Fishkin, 2024; 2026). When Google shows an AI-generated summary, users click a conventional result on 8 percent of visits, roughly half the 15 percent observed when no summary appears, and only 1 percent click a source cited inside the summary itself (Pew Research Center, 2025). ChatGPT reported on the order of 900 million weekly active users by early 2026, and industry trackers estimate that roughly a third of consumers now begin information and product journeys inside an AI assistant rather than a search box (Demandsage, 2026). The search results page, the internet's front door for a quarter century, is being replaced as the first point of touch. What is dying is not search itself; people ask more of the internet than ever. What is dying is the search engine as a destination, the page of ranked links that turned questions into traffic.

This paper examines that replacement with three questions in mind. First, how did we get here: what is the actual sequence of changes in how people first meet online content, and what mechanism drove each transition? Second, how does intent behave in the new environment: which of the classic query intents survive the arrival of synthesized answers, which are absorbed, and what does the shape of expressed demand look like when a query can be a full sentence with constraints instead of two keywords? Third, what follows for the organizations that depend on being discovered: if the click is no longer the unit of value, what is, and how should firms allocate effort?

The paper makes four contributions. It assembles the historical progression of content discovery into a five-era framework grounded in the information retrieval literature. It offers a compact formal treatment that connects the ranking mathematics of the SERP era, the attention mechanism that made synthesis possible, and an expected-visit model that shows precisely where the economics change. It reviews the strongest available behavioral evidence, from clickstream panels and controlled click-through studies to conversion data on AI-referred visitors, and it distills that evidence into a set of concrete, testable imperatives for practitioners. Throughout, the aim is to keep the technical material in service of the argument rather than the reverse.

A note on method. This is a synthesis and modeling paper, not a primary empirical study. The behavioral data cited come from industry measurement firms and one large research nonprofit; their

limitations are discussed in Section 8. Peer-reviewed sources anchor the technical and taxonomic claims. Where industry figures are used, the paper prefers studies that disclose sample construction and method.

## 2. Background and Related Work

Three research streams meet in this paper: classical information retrieval, the study of search intent, and the neural architecture work that produced modern language models.

Classical retrieval treats discovery as a matching problem between a short query and a large corpus. Term-weighting schemes such as TF-IDF formalized the intuition that words which are frequent in a document but rare in the corpus carry the most signal (Salton and Buckley, 1988). The probabilistic relevance framework refined this into BM25, still a workhorse of lexical retrieval (Robertson and Zaragoza, 2009). Link analysis added a second, query-independent signal: PageRank modeled a page's authority as the stationary probability that a random surfer lands on it, using the web's link graph as a vast implicit voting system (Brin and Page, 1998; Page et al., 1999).

The intent literature begins with Broder (2002), who observed that the need behind a web query is often not informational at all. His taxonomy separates informational queries (the user wants to know something), navigational queries (the user wants to reach a particular site), and transactional queries (the user wants to perform an action such as purchasing or downloading). Jansen, Booth, and Spink (2008) operationalized the taxonomy at scale and found informational needs dominant, at roughly four-fifths of queries, with navigational and transactional needs splitting the remainder. Eye-tracking and log studies established how strongly attention and clicks concentrate at the top of a ranked list, and how much of that concentration is positional trust rather than assessed relevance (Granka, Joachims, and Gay, 2004; Joachims et al., 2005). This position bias became the load-bearing assumption of an entire industry: search engine optimization exists because rank determines traffic.

The third stream made answers possible. Vaswani et al. (2017), in the paper titled 'Attention Is All You Need,' replaced recurrent sequence processing with self-attention, an operation that lets every token in a sequence weigh its relationship to every other token in parallel. The Transformer architecture that resulted trained faster, scaled further, and transferred better than its predecessors. Its descendants moved into search in stages: BERT improved Google's interpretation of natural-language queries from 2019 (Devlin et al., 2019; Nayak, 2019), dense passage retrieval replaced some lexical matching with learned embeddings (Karpukhin et al., 2020), and retrieval-augmented generation coupled a retriever to a generator so that a model could compose fluent text grounded in fetched documents (Lewis et al., 2020). Aggarwal et al. (2024) supplied the first systematic study of visibility inside such systems, introducing generative engine optimization and showing that content-side interventions can raise a source's visibility in generated answers by up to 40 percent.

A fourth, younger stream measures what all of this does to behavior: clickstream analyses of zero-click search (Fishkin, 2024; 2026), the Pew Research Center's browsing-panel study of AI summaries (2025), Ahrefs' large-sample click-through comparisons (Soulo, 2025), and Adobe's traffic and conversion telemetry on AI-referred visitors (Adobe, 2025; 2026). Section 5 examines this evidence in detail.

## 3. The Five Eras of First Touch

Each era of content discovery is defined by who or what performs the work of selection, and each transition moved that work further away from the user. Figure 1 places the eras on a timeline. The dates overlap because eras fade rather than end; directories persisted for years after crawlers arrived, and the ranked list still serves billions of queries today.

**Figure 1. Five eras of first touch: where content discovery begins, 1990 to present**

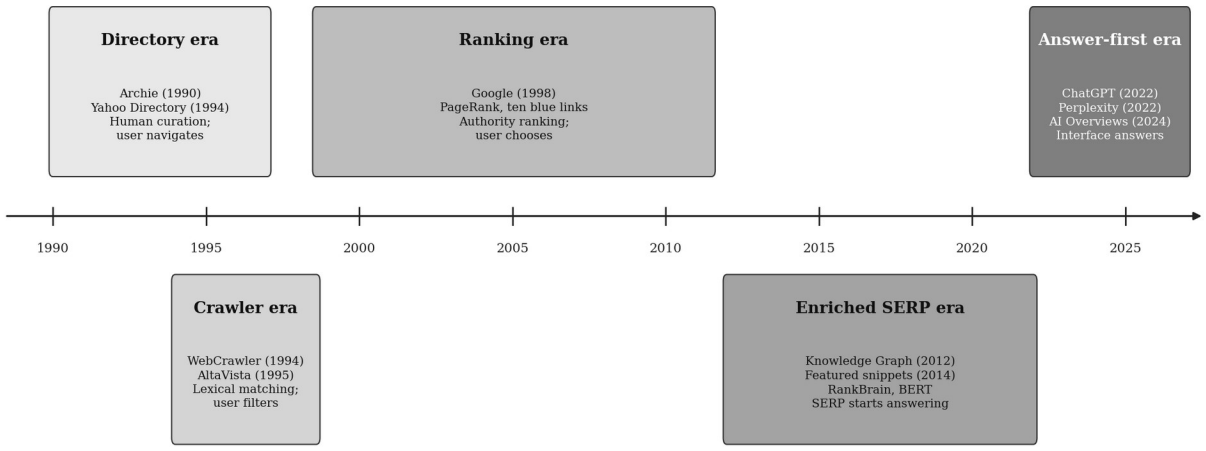


*Figure 1. Five eras of first touch: where content discovery begins, 1990 to present. Eras overlap because each fades rather than ends.*

### *3.1 The directory era (1990 to 1997): the user navigates*

The earliest discovery tools indexed almost nothing. Archie, built by Alan Emtage in 1990, searched only the file names on FTP servers. The Yahoo Directory, launched by Jerry Yang and David Filo in 1994, was a human-curated catalog organized into browsable categories rather than a search engine in the modern sense. Selection was editorial and navigation was manual: the user descended a category tree and did the finding herself. For content owners, discovery meant a listing, and the scarce resource was a place in the catalog.

### *3.2 The crawler era (1994 to 1998): the user filters*

Full-text crawling changed the unit of discovery from the site to the page. WebCrawler brought full-text search in 1994; AltaVista, launched in 1995, handled around 13 million queries per day at its peak and made keyword search a mainstream habit. Matching was lexical, which meant results were only as good as the overlap between the user's words and the page's words. Ranking was primitive, so the burden of filtering fell on the user, who paged through results and learned, by necessity, to craft queries. The era's business pathology followed directly: keyword stuffing worked because term frequency was the signal.

### *3.3 The ranking era (1998 to 2011): the user chooses*

Google's insight was that the web itself contained a relevance signal no page could easily fake: the link. PageRank ranked pages by the structure of endorsements around them, and the combination of link authority with lexical relevance produced results good enough that scanning ten of them usually sufficed (Brin and Page, 1998). The ranked list became the standard interface of the internet, and an entire economy formed around position on it. Search engine optimization, paid search, and the content-publishing boom of the 2000s all rest on the same behavioral fact documented by the eye-tracking studies: attention decays steeply with rank, so rank is revenue. In this era the engine selected and ordered, but the user still chose, clicked, and read on the destination site. Value reached publishers through the click.

### *3.4 The enriched-SERP era (2012 to 2022): the page starts answering*

The results page then began to absorb the answer. Google's Knowledge Graph, announced in May 2012, let the engine present facts about entities directly. Featured snippets, appearing from 2014, extracted a passage from a ranked page and displayed it above the list. Under the surface, the Hummingbird rewrite (2013), RankBrain (2015), and BERT (2019) moved interpretation from keywords toward meaning, so that conversational queries began to work. Each of these features traded a click for convenience. The zero-click search, a query resolved entirely on the results page, became a measurable and growing category during this period, and the tension it created with publishers foreshadowed the current era: the engine was beginning to compete with its own sources for the user's attention.

### *3.5 The answer-first era (2022 to present): the interface answers*

ChatGPT's launch in November 2022 demonstrated that a language model could serve as a general question-answering interface, and Perplexity, launched weeks later, showed the same interface grounded in live retrieval with citations. Google responded by moving synthesis into search itself: the Search Generative Experience in 2023, AI Overviews rolled out broadly from May 2024, and a fully conversational AI Mode thereafter. The structural change is larger than any single product. In the previous eras the system retrieved and ordered documents while the human extracted the answer. Now the system extracts and composes the answer, and the documents become raw material. The ranked list has not disappeared, but it has been demoted from the interface to a fallback, and the first point of touch, for a growing share of users and queries, is a rendered answer. No rival ever beat the ten blue links. They lost to a paragraph.

## 4. A Formal Sketch of Discovery

A note before the mathematics. This section formalizes the progression described so far and is written mainly for readers who build or study these systems. Everyone else can skip ahead to Section 5, which covers the behavioral evidence, without losing the thread.

A small amount of formalism makes the transition precise and shows exactly where the economics shift. This section presents the scoring logic of each era, the attention mechanism that enabled synthesis, and an expected-visit model that decomposes a source's traffic under the answer layer. Figure 2 compares the two pipelines end to end.

**Figure 2. Two pipelines of discovery: ranked retrieval versus retrieve-then-synthesize**

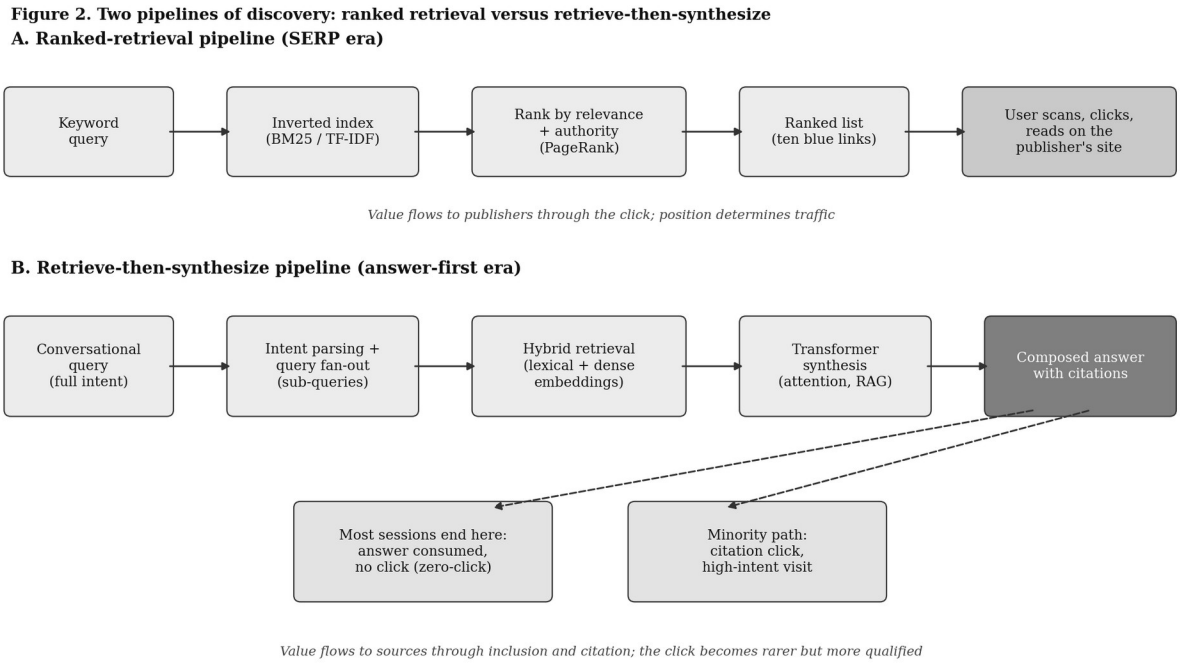


*Figure 2. Two pipelines of discovery. Panel A: ranked retrieval, where value flows to publishers through the click. Panel B: retrieve-then-synthesize, where value flows to sources through inclusion and citation.*

### *4.1 Lexical relevance and link authority*

The crawler and ranking eras scored documents with functions a person can inspect. BM25 scores a document d against a query q as a sum over query terms, weighting each term's in-document frequency f(t, d) against its corpus rarity IDF(t), with $k_1$ and b as saturation and length-normalization constants (Robertson and Zaragoza, 2009):

$$score(q,\ d) = \sum_{t \in q} IDF(t) \cdot \frac{f(t,\ d) \cdot (k_1 + 1)}{f(t,\ d) + k_1 \left(1 - b + b \cdot \frac{|d|}{avgdl}\right)} \qquad (1)$$

PageRank added the query-independent authority signal, defining a page's score recursively over the set $B_p$ of pages linking to it, where L(q) is the number of outbound links on page q, N the number of pages, and λ a damping factor near 0.85 (Brin and Page, 1998):

$$PR(p) = \frac{1 - \lambda}{N} + \lambda \sum_{q \in B} \frac{PR(q)}{L(q)} \qquad (2)$$

Both equations rank existing documents. Neither can produce a sentence. The move from Equation 2 to a composed answer required an architecture that could model language itself.

### *4.2 Attention and the transformer*

Vaswani et al. (2017) built sequence models from a single primitive, scaled dot-product attention. Queries Q, keys K, and values V are learned projections of the token sequence; each token's output is a weighted mixture of all values, with weights given by the softmax of scaled query-key similarities:

$$Attention(Q,\ K,\ V) = softmax\left(\frac{Q K^T}{\sqrt{d_k}}\right)V \qquad (3)$$

Two properties of Equation 3 matter for this paper's story. Attention is parallel, which made training on web-scale text economically feasible, and it is content-addressed, meaning every token can draw on any other token regardless of distance. Language models built by stacking this operation learned to interpret queries, judge relevance, and generate prose. The same primitive therefore serves three roles in the modern discovery stack: understanding the question, matching it to passages, and writing the answer. It is the reason the eras described in Section 3 converged on synthesis rather than better ranking.

### *4.3 Dense retrieval and retrieval-augmented generation*

Transformer encoders also changed retrieval itself. In the bi-encoder design of dense passage retrieval, a query encoder $E_Q$ and a passage encoder $E_P$ map text into a shared vector space, and relevance becomes an inner product (Karpukhin et al., 2020):

$$\mathrm{sim}(q,\ p) = E_Q(q)^T E_P(p) \qquad (4)$$

Dense retrieval matches meaning rather than surface strings: a passage can rank highly for a query that shares none of its words. Retrieval-augmented generation then closes the loop. The system retrieves a small set $Z_k(x)$ of passages for input x, and generates output y by marginalizing over them, with $p_\eta$ the retriever's relevance distribution and $p_\theta$ the generator (Lewis et al., 2020):

$$P_{RAG}(y \mid x) \approx \sum_{z \in Z_k(x)} p_\eta(z \mid x) \cdot p_\theta(y \mid x,\ z) \qquad (5)$$

Equation 5 is the answer engine in miniature, and it carries a consequential asymmetry. A source influences the answer only if it enters $Z_k(x)$, the retrieved set, and it earns attribution only if the interface surfaces it as a citation. Presence in the index is no longer sufficient; presence in the retrieved-and-cited set is the new threshold of existence.

### *4.4 An expected-visit model*

Consider a source s (a site, a brand, a publisher) and a universe of queries Q, where query q occurs with volume $\omega_q$. Let $\alpha_q$ be the share of impressions for q that resolve through the ranked list, $r_s(q)$ the source's rank when it appears, $\pi_s(q)$ the probability that the answer layer includes and cites s for query q, and $c_s(q)$ the probability of a click on that citation. Expected visits decompose as:

$$E[V_s] = \sum_{q \in Q} \omega_q \left[\alpha_q \cdot CTR(r_s(q)) + (1 - \alpha_q) \cdot \pi_s(q) \cdot c_s(q)\right] \qquad (6)$$

Click-through on ranked lists decays steeply and approximately as a power law in rank, with decay exponent γ greater than one:

$$CTR(r) \approx \kappa \cdot r^{-\gamma}, \quad \gamma > 1 \qquad (7)$$

The entire transition of Section 3.5 lives inside two movements of Equation 6. First, $\alpha_q$ is falling: a growing fraction of impressions is resolved by the answer layer rather than the list, and the evidence in Section 5 suggests the fall is fastest for informational queries. Second, within the answer term, $\pi_s$ replaces rank as the strategic variable. Rank is ordinal, public, and contested inch by inch; inclusion is probabilistic, opaque, and winner-take-most, since a synthesized answer typically cites a handful of sources drawn from a candidate pool of thousands. The practical program of Section 7 is, in these terms, a program for raising $\pi_s$ and $c_s$ while defending the shrinking but still substantial $\alpha_q$ term. Note also what the model implies about measurement: a source can gain influence (appearing in answers, shaping decisions) while its measured visits fall, because the brand enters the user's awareness without producing a session. Influence and traffic have come apart, and firms that measure only traffic will misread their own position.

## 5. What the Behavioral Evidence Shows

Four bodies of evidence, produced with different methods by unrelated organizations, point the same direction. Table 1 summarizes them.

*Table 1. Selected evidence on user behavior in the answer-first era*

| Source | Sample and method | Key finding |
|---|---|---|
| Datos clickstream panel (SparkToro), 2024 | Panel of millions of US and EU devices | 58.5% of US Google searches ended without a click to the open web; 360 of every 1,000 searches produced an open-web click |
| Datos clickstream panel (SparkToro), 2026 | Same panel, January to April 2026 | 68.0% of US Google searches ended without a click; open-web clicks fell to 276 per 1,000 searches |
| Pew Research Center, 2025 | Browsing records of 900 US adults, March 2025 | Users clicked a result on 8% of searches that showed an AI summary versus 15% without one; only 1% clicked a source cited inside the summary; 26% of AI-summary pages ended the session |
| Ahrefs, 2025 | 300,000 keywords, half with AI Overviews | Presence of an AI Overview correlated with a 34.5% lower click-through rate for the top-ranking page in April 2025, widening to 58% by December 2025 |
| Adobe Digital Insights, 2025 to 2026 | Aggregated visits to US retail and travel sites | AI-referred traffic grew roughly tenfold between July 2024 and February 2025; by May 2026 it converted 54% better than non-AI traffic, with lower bounce rates and longer visits |

### *5.1 The zero-click trajectory*

The Datos clickstream panel analyzed by SparkToro provides the longest consistent series. In 2024, 58.5 percent of US Google searches ended without a click to the open web; for every 1,000 searches, 360 clicks reached open-web sites, with the remainder absorbed by answers on the page, refinements, abandonment, or clicks back into Google's own properties (Fishkin, 2024). By the first four months of 2026

the zero-click share had reached 68 percent and open-web clicks had fallen to 276 per 1,000 searches (Fishkin, 2026). The direction is not new, since zero-click behavior grew throughout the enriched-SERP era, but the slope steepened after AI Overviews reached broad deployment, appearing on more than a fifth of queries in the same period.

### *5.2 Behavior in the presence of AI summaries*

The Pew Research Center study is the cleanest available comparison of the same users with and without the answer layer. Among 900 US adults who shared browsing records in March 2025, searches that returned an AI summary produced a click on a conventional result 8 percent of the time, against 15 percent when no summary appeared. Clicks on sources cited inside the summary occurred on 1 percent of such pages. Users also ended their browsing session outright more often after seeing a summary, 26 percent versus 16 percent (Pew Research Center, 2025). Read against Equation 6, this is a direct observation of $\alpha$ falling and of $c$, the citation click-through, sitting at a level that makes inclusion valuable mainly for influence rather than referral volume.

Ahrefs' keyword-matched comparison sharpens the picture for publishers. Across 300,000 informational keywords, half with AI Overviews present, the top-ranking page's click-through rate was 34.5 percent lower under an AI Overview in April 2025, and the gap had widened to 58 percent by December 2025 (Soulo, 2025). The penalty falls precisely on the position that the entire SEO industry spent two decades competing for. These are correlational designs, and query mix and seasonality complicate attribution, a caveat taken up in Section 8; but the magnitude, replication across datasets, and coherence with the panel data leave little doubt about the direction.

### *5.3 The compensating signal: quality of the surviving click*

The same period shows the opposite movement in traffic quality. Adobe's telemetry across US retail and travel sites recorded a roughly tenfold increase in AI-referred traffic between July 2024 and February 2025, and the visitors behaved differently on arrival: bounce rates 23 percent lower than other channels, more pages per visit, and substantially longer sessions (Adobe, 2025). By March 2026, AI-referred retail traffic converted 42 percent better than non-AI traffic, a full reversal from a year earlier when it converted worse, and by May 2026 the advantage stood at 54 percent (Adobe, 2026).

The two movements are one phenomenon. The answer layer absorbs the early, exploratory stages of a journey: definitions, comparisons, suitability checks. A user who clicks through after that absorption has already been qualified by the conversation, arriving later in the funnel with narrower intent. Discovery traffic is not so much disappearing as being refined: fewer visits, further along, worth more each. For businesses this reframes the loss of raw traffic as a compression of the funnel, with the top of it now happening on infrastructure the business does not control and mostly cannot observe.

## 6. How Intent Works Now

Broder's taxonomy was built for an interface that could only return documents, so every intent had to be expressed as a fetch. Synthesis relaxes that constraint, and the three classes respond differently. Figure 3 traces the routing.

**Figure 3. Intent routing in the answer-first stack: what synthesis absorbs and what it leaves intact**

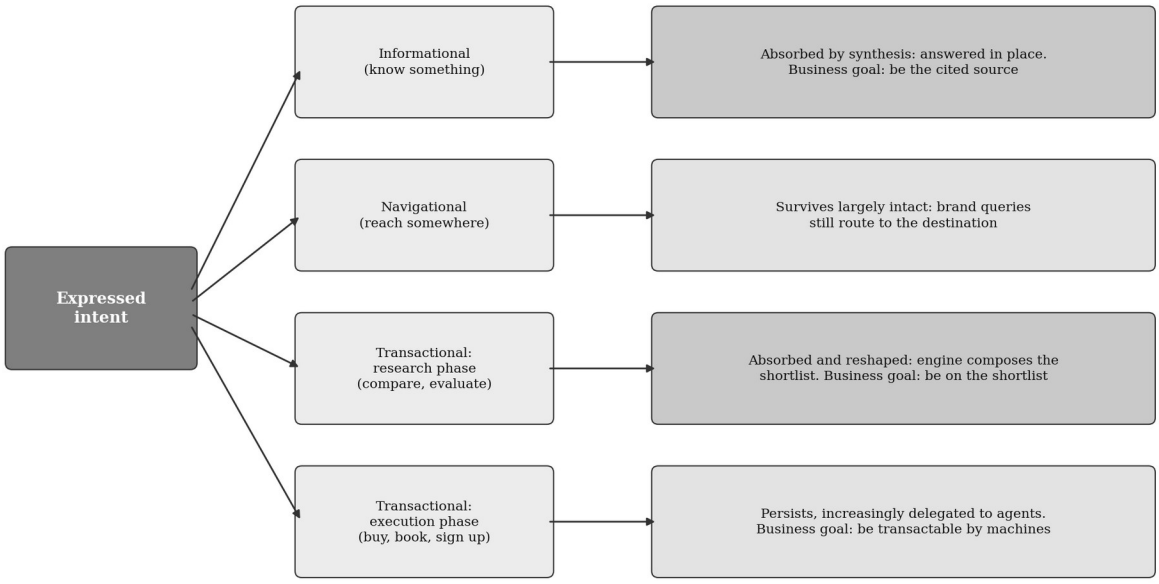


*Intent classes follow Broder (2002); resolution paths reflect observed behavior in answer engines (Sections 5 and 6).*

*Figure 3. Intent routing in the answer-first stack. Intent classes follow Broder (2002); resolution paths reflect the behavioral evidence of Section 5.*

Informational intent, the majority class in every measurement since Jansen et al. (2008), is the natural prey of synthesis. A question whose answer is a fact, an explanation, or a comparison is resolved more efficiently by a composed paragraph than by a reading list, and the zero-click data of Section 5 show users accepting that trade at scale. For content built to intercept informational queries, the implication is stark: the query still exists, but its resolution no longer requires a visit. What remains contestable is authorship of the answer, that is, being the source the engine retrieves, trusts, and cites.

Navigational intent survives nearly intact. A user who wants a specific destination is not asking for prose, and engines still route brand and site queries directly. This has a strategic corollary that is easy to miss: as generic informational traffic is absorbed, demand that arrives as your name becomes a larger share of what reaches you, and building memorable, nameable brands quietly becomes a discovery strategy rather than only a marketing one.

Transactional intent splits under pressure. Its research phase, the comparing and shortlisting that used to generate dozens of queries and site visits, is absorbed into the conversation: the engine now composes the shortlist. Its execution phase, the purchase or booking or signup, still requires the destination, and it increasingly arrives through agents acting on the user's behalf. The business goal shifts accordingly, from

capturing research traffic to being present on machine-composed shortlists and being transactable when an agent, not a person, is the one completing the task.

The expression of intent changes along with its routing. Keyword queries were a compression format imposed by the interface: a user with a rich need learned to strip it to two or three tokens and iterate. Conversational interfaces remove the compression, and observed queries in answer engines run several times longer than search queries, carrying constraints, context, and preferences in a single utterance. The engine responds with query fan-out, decomposing the utterance into sub-queries it retrieves against in parallel. One consequence is session compression: journeys that took eight searches take one or two turns. Another is that intent data, the exhaust that made search advertising the most profitable business model of the modern internet, now accumulates inside conversational platforms at far higher resolution than any keyword log, with consequences for competition and privacy that research is only beginning to address.

## 7. What Businesses Need to Build

Equation 6 organizes the practical agenda. A firm's discoverability now rests on five quantities: the surviving ranked-list term (rank r and its CTR), the inclusion probability π, the citation click-through c, and the conversion value of each visit. The imperatives below map onto those quantities. None of them replaces the fundamentals of being genuinely useful; all of them are ways of making that usefulness legible to machines.

### *7.1 Be retrievable*

Nothing else matters if the retriever cannot reach, parse, and segment your content. That begins with unglamorous infrastructure: allowing reputable AI crawlers in robots.txt, serving fast and stable pages, and rendering substantive content in HTML rather than behind client-side scripts. It extends to machine readability: structured data in JSON-LD (Organization, Product, FAQ, and Article schemas) gives parsers unambiguous facts, and emerging conventions such as llms.txt offer a curated map of a site's most important content for language-model consumption. Retrieval systems operate on passages, not pages, so content should be written in self-contained blocks: a heading that states the question, a passage that answers it completely within a few sentences, and supporting detail after. A passage that requires the surrounding page to be intelligible is a passage the synthesis layer cannot use.

### *7.2 Be citable*

Aggarwal et al. (2024) provide the only controlled evidence to date on what raises visibility inside generated answers, and their findings are conveniently concrete: adding citations to authoritative sources, quotations from identifiable experts, and specific statistics raised a source's visibility in answers by up to 40 percent, while keyword stuffing, the legacy tactic, did nothing or worse. The pattern is intuitive once Equation 5 is in view. A generator assembling a grounded answer favors passages that carry verifiable, attributable, quotable material, because such passages lower its own cost of being right. Original data, named authorship, precise claims, and clear sourcing are no longer merely good editorial practice; they are the selection criteria of the new gatekeeper. Lower-ranked sources have the most to gain, since the answer layer draws from a wider pool than the old top three positions.

### *7.3 Be corroborated*

Answer engines learn about a brand from the entire retrievable web, not from the brand's own site alone. When a user asks for the best option in a category, the retrieved set is dominated by third-party surfaces: review platforms, community threads, comparison articles, industry publications. A firm absent from those surfaces is absent from the candidate pool no matter how good its own content is. Corroboration work therefore looks less like link building and more like reputation building: earning presence in the specific pages engines already retrieve and cite for the queries that matter, keeping factual details (name, category, pricing, positioning) consistent everywhere they appear so that entity-resolution systems consolidate rather than fragment the brand, and treating communities such as Reddit, which answer engines cite heavily, as places to be genuinely present rather than to astroturf.

### *7.4 Be measurable*

The answer-first funnel is invisible to the instruments most firms run. A server log records the visit, but the impression that mattered, the brand's appearance inside a rendered answer, happens on infrastructure that emits no signal to the source. Measurement therefore has to be reconstructed from the outside: sampling the major engines with representative prompts and recording whether the brand appears, where it ranks within the answer, which sources are cited, and with what sentiment. A new instrumentation category has grown around exactly this need, and purpose-built platforms such as AnswerRank (answerrank.so) track a brand's presence, position, and sentiment across AI-generated answers over time and translate the gaps into prioritized fixes. Whatever tooling a firm chooses, the metrics that matter map directly onto the model of Section 4: inclusion rate ($\pi$) by query and engine, citation share against competitors, and the conversion rate of AI-referred sessions, which analytics platforms still routinely misattribute to direct traffic. What is not measured here will be misread as decline.

### *7.5 Be worth the click, and ready for its buyer*

The clicks that survive synthesis are fewer and better, and the destination should be built for them. A visitor arriving from an answer has already consumed the overview; landing her on a generic explainer wastes the qualification the conversation performed. Deep, specific landing content, transparent pricing, and low-friction next steps convert this traffic, and Adobe's conversion figures suggest the payoff is already material (Adobe, 2026). Two longer-horizon investments follow from Section 6. First, first-party relationships (subscriptions, communities, direct channels) hedge against an intermediary that no longer forwards attention reliably. Second, agent-readiness: as execution-phase transactions migrate to software agents, clean product feeds, documented APIs, and checkout flows that a machine can traverse become distribution assets in their own right. The firms that treated the mobile transition as a checkbox rediscovered its cost for a decade; the agentic transition deserves earlier seriousness.

## 8. Limitations and Future Research

The evidence base has real weaknesses. Clickstream panels overrepresent desktop and consenting users; the Pew sample, though carefully constructed, covers a single month of one country's behavior; and the click-through comparisons are correlational, since Google does not randomize the presence of AI

Overviews for external observers. Query mix shifts confound before-and-after comparisons: if AI Overviews appear disproportionately on queries that were already low-click, naive estimates overstate the causal penalty. Industry sources, however transparent about method, are not peer reviewed, and several have commercial stakes in the phenomena they measure. The formal model of Section 4 is deliberately spare; it treats inclusion as a static probability when in practice $\pi$ is the output of adversarial, constantly retrained systems, and it ignores the strategic interaction between engines and publishers, including licensing, paywalls, and blocking, that will shape the equilibrium.

Each weakness is a research opening. The field needs standardized, auditable measures of citation share and inclusion rate across engines; causal designs, perhaps exploiting the staggered rollout of AI features across countries and query classes; welfare analysis of the publisher ecosystem under falling referral traffic, since the answer layer consumes a commons it does not directly replenish; and study of agent-mediated transactions, where the buyer of record is software and the entire apparatus of persuasion built for human eyes must be rethought. The intent taxonomy itself deserves renewal: a conversational turn carrying five constraints and an implicit budget is poorly described by any of Broder's three boxes.

## 9. Conclusion

Thirty years of content discovery compress into a single sentence: the work of finding moved from the user to the machine, one stage at a time, until the machine began delivering not the location of an answer but the answer itself. Directories asked the user to navigate, crawlers to filter, ranked lists to choose; enriched result pages started answering the easy questions, and transformer-based engines now answer the hard ones, drawing on the open web as raw material and citing a favored few of its sources. Each transition rewarded the businesses that understood it early, and this one is no different in kind, only in speed.

The evidence assembled here supports a plain reading. The click is becoming scarce: two-thirds of searches already end without one, and the presence of a synthesized answer cuts the survivors sharply. The click is simultaneously becoming precious: visitors who arrive through an answer convert at rates their keyword-era predecessors never approached. Between those two facts sits the new strategic terrain, where the contested variable is no longer position on a page but probability of inclusion in an answer, and where influence and traffic no longer travel together. Businesses that learn to be retrievable, citable, corroborated, measurable, and worth the smaller stream of clicks will find the answer-first web navigable. Those that keep optimizing for a results page their customers no longer see will be measuring their own disappearance with instruments built for a previous internet. The questions are not going anywhere. The place we took them is.

## References


Adobe. (2025). *The explosive rise of generative AI referral traffic.* Adobe Digital Insights, Adobe Business Blog. https://business.adobe.com/blog/the-explosive-rise-of-generative-ai-referral-traffic

Adobe. (2026). *AI-driven traffic surges across industries, retail sees biggest gains.* Adobe Digital Insights, Adobe Business Blog. https://business.adobe.com/blog/ai-driven-traffic-surges-across-industries

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., and Deshpande, A. (2024). GEO: Generative Engine Optimization. *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, 5 to 16. https://doi.org/10.1145/3637528.3671900

Brin, S., and Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1 to 7), 107 to 117.

Broder, A. (2002). A taxonomy of web search. *ACM SIGIR Forum*, 36(2), 3 to 10.

Demandsage. (2026). *ChatGPT statistics: Latest active user data.* https://www.demandsage.com/chatgpt-statistics/

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019*, 4171 to 4186.

Fishkin, R. (2024). *2024 zero-click search study: For every 1,000 US Google searches, only 374 clicks go to the open web.* SparkToro, with Datos. https://sparktoro.com/blog/2024-zero-click-search-study-for-every-1000-us-google-searches-only-374-clicks-go-to-the-open-web-in-the-eu-its-360/

Fishkin, R. (2026). *In 2026, less than one third of Google searches still send a click.* SparkToro, with Datos. https://sparktoro.com/blog/in-2026-less-than-one-third-of-google-searches-still-send-a-click/

Granka, L. A., Joachims, T., and Gay, G. (2004). Eye-tracking analysis of user behavior in WWW search. *Proceedings of the 27th Annual International ACM SIGIR Conference*, 478 to 479.

Jansen, B. J., Booth, D. L., and Spink, A. (2008). Determining the informational, navigational, and transactional intent of Web queries. *Information Processing and Management*, 44(3), 1251 to 1266.

Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. (2005). Accurately interpreting clickthrough data as implicit feedback. *Proceedings of the 28th Annual International ACM SIGIR Conference*, 154 to 161.

Karpukhin, V., Oğuz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., and Yih, W. (2020). Dense passage retrieval for open-domain question answering. *Proceedings of EMNLP 2020*, 6769 to 6781.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., and Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459 to 9474.

Nayak, P. (2019). *Understanding searches better than ever before.* Google, The Keyword blog. https://blog.google/products/search/search-language-understanding-bert/

Page, L., Brin, S., Motwani, R., and Winograd, T. (1999). *The PageRank citation ranking: Bringing order to the web.* Technical report, Stanford InfoLab.

Pew Research Center. (2025). *Google users are less likely to click on links when an AI summary appears in the results.* https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/

Robertson, S., and Zaragoza, H. (2009). The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4), 333 to 389.

Salton, G., and Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information Processing and Management*, 24(5), 513 to 523.

Soulo, T. (2025). *AI Overviews reduce clicks by 34.5%*, and update: *AI Overviews reduce clicks by 58%*. Ahrefs Blog. https://ahrefs.com/blog/ai-overviews-reduce-clicks/ and https://ahrefs.com/blog/ai-overviews-reduce-clicks-update/

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998 to 6008.